arXiv:1402.1956v1 [cs.AI] 9 Feb 2014

Revisiting the Learned Clauses Database Reduction Strategies

Said Jabbour and Jerry Lonlac and Lakhdar Sais and Yakoub Salhi

CRIL - CNRS, Université d'Artois, Lens, France
{jabbour, lonlac, sais, salhi}@cril.fr

Abstract. In this paper, we revisit an important issue of CDCL-based SAT solvers, namely the learned clauses database management policies. Our motivation takes its source from a simple observation on the remarkable performances of both random and size-bounded reduction strategies. We first derive a simple reduction strategy, called Size-Bounded Randomized strategy (in short SBR), that combines maintaing short clauses (of size bounded by k), while deleting randomly clauses of size greater than k. The resulting strategy outperform the state-of-the-art, namely the LBD based one, on SAT instances taken from the last SAT competition. Reinforced by the interest of keeping short clauses, we propose several new dynamic variants, and we discuss their performances.

1 Introduction

Today, SAT has gained a considerable audience with the advent of a new generation of solvers able to solve large instances encoding real-world problems. These solvers, often called *modern SAT solvers* [22,11], are based on the classical DPLL procedure [9] enhanced with: (i) an efficient implementation of unit propagation through incremental and lazy data structures, (ii) restart policies [15,20], (iii) activity-based variable selection heuristics (VSIDS-like) [22], and (iv) *clause learning* [27,22]. Clause learning is now recognized as one of the most important component of Modern SAT solvers. The main idea is that when a current branch of the search tree leads to a conflict, clause learning aims to derive a clause that succinctly expresses the causes of the conflict. Such learned clause is then used to prune the search space. Clause learning also known in the literature as Conflict Driven Clause Learning (CDCL) refers now to the most known and used First UIP learning scheme, first integrated in the SAT solver Grasp [27] and efficiently implemented in zChaff [22]. Most of the SAT solvers, integrate this strong learning scheme. Theoretically, by integrating clause learning to DPLL-like procedures [9], the obtained SAT solver formulated as a proof system is as powerful as general resolution [24,26].

In practice, the efficiency of CDCL-based SAT solvers heavily depends on the strategy used to manage the learned clauses database. Indeed, as at each conflict, a new clause is added to the learned clauses database, its size grows exponentially. To avoid this combinatorial explosion, several learned clauses database

management strategies have been proposed (e.g. [22,11,1,3,16]). These strategies aims to maintain a learned clause database of reasonable size by eliminating clauses deemed irrelevant to the subsequent search. All these management strategies follow a predefined cleaning time-sequence where the interval between two successive reduction steps is more or less important. At each conflict, an activity is associated to the learned clause (static strategy). Such heuristic-based activity aims to weight each clause according to its relevance to the search process. In the case of dynamic strategies, such clauses activities are dynamically updated. The reduction of the learned clauses database consists in eliminating inactive or irrelevant clauses. Even if all the proposed learned clause deletion strategies are shown to be empirically efficient, determining the most relevant clause to the search process remains a challenging task. It is important to note that the efficiency of most of the state-of-the-art learned clauses management strategies heavily depends on the cleaning frequency and on the amount of clauses to be deleted each time.

Different implementations of SAT solvers are proposed each year to the SAT competition, they include several enhancements to the main components of CDCL-based solvers. The SAT competitions, races and challenges stimulate the development of efficient implementations. However, such a race to the most efficient implementation has led to increasingly complex solvers with several number of parameters. Such parameters are either static (fixed before search), or dynamic, their values are conditionally set during search by observing the search behavior. These sophisticated implementations increase the difficulty to understand what is essential from what is not. In [19], an empirical analysis focusing on the principal techniques that have contributed to the impressive performance of modern SAT solvers has been conducted. This is really a first step towards a deep understanding of modern SAT solvers.

In this paper, we revisit an important issue of CDCL-based SAT solvers, namely the learned clauses database management policy. It is *important to note* at this point that considering short clauses as the most relevant ones (size-bounded based reduction strategies) have been proposed since 1996 by Marques Silva and Sakallah (Grasp [27]) and Bayardo and Schrag (RelSAT [6]). Most of the state-of-the-art SAT solvers maintain systematically only binary clauses, while for clauses of size greater that two, several sophisticated relevance based measures have been proposed to predict the most relevant ones.

Our motivation in this work has its source from a simple observation on the remarkable performances of size-bounded reduction strategies. From this first results, we decided to investigate other "naive" strategies such as random and FIFO (First In First Out). The goal is to quantify the performance gap between these simple strategies and those implemented in the state-of-the-art SAT solvers. Secondly, we derive a simple reduction strategy, called Size-Bounded Randomized Strategy (in short SBR), that combines maintaing short clauses (of size bounded by k), while deleting randomly clauses of size greater than k. Reinforced by the interest of keeping short clauses, we propose several new dynamic measures that allow us to quantify the relevance of a given learned clause w.r.t.

the current state of the search process. Intuitively, *a short learned clause with literals assigned most often at the top of the current branch of the search tree is considered more relevant*. Several relevance-based strategies are then derived, allowing us to keep learned clauses that are more likely to cut branches at the top of the search tree. All these strategies are integrated at the top of MiniSAT 2.2 and compared with the state-of-the-art SAT solvers on application instances taken from the last SAT competition. To confirm the superiority of size-based measures against LBD-based ones, we also present the results obtained by substituting LBD with dynamic size-based activity inside *Glucose* 3.0. The obtained results bring up to date old strategies, such as size-bounded proposed more than fifteenth years ago [27,5,6] (see Section 2 - Related Works). These results are not surprising, since short clauses are usually preferred for their ability to reduce further the search space. This is also why all the SAT solvers keep the learned binary clauses.

The paper is organized as follows. We first recall related works in Section 2. Then, after some technical background and preliminary definitions (Section 3), we motivate our study by investigating the size-bounded, random and FIFO deletion strategies in Section 4. In Section 5, we present our first reduction strategy, called Size-Bounded Randomized strategy. In Section 6, we describe several relevant based deletion strategies, that allow us to keep learned clauses that are more likely to cut branches at the top of the search tree. All these new strategies are integrated at the top of MiniSAT 2.2 and compared to the state-of-the-art SAT solvers *Glucose* 3.0 and *Lingeling* 2013. We also present the results obtained by integrating size-bounded based strategies inside *Glucose* 3.0. A discussion is provided in Section 7, before concluding.

2 Related Works

In this section, we give a non exhaustive review of some important deletion strategies as described by the authors. To our knowledge, maintaing a relevant set of clauses goes back to the developments of efficient resolution-based proof procedures in automated theorem proving. In [12], D. Gelperin proposed different strategies that attempt to determine the satisfiability of a set of input clauses while at the same time minimizing the cardinality of the set of retained clauses.

In constraint satisfaction and propositional satisfiability problems, to overcome the overhead of unrestricted learning, several strategies have been proposed by Bayardo et al. [6,5], including size-bounded and relevance-bounded learnings. They defined size-bounded (respectively, relevance-bounded) learning of order i, which retains indefinitely only clauses derived reasons containing i or fewer variables (respectively, maintains any reason that contains at most i variables whose assignments have changed since the reason was derived). In SAT, this issue is first investigated in GRASP by Joao Marques Silva et al. [27]. Indeed, in order to guarantee the worst case growth of the clause database to be polynomial in the number of variables, in [27] the authors propose a selective strategy on the clauses that have to be added to the clause database. More precisely, given an

integer parameter k, conflict-induced clauses whose size (number of literals) is less than k are marked green (added to the database) while those of size greater than k are marked red and kept around only while they are unit clauses, i.e., a red clause is deleted when it contains more than one free (unassigned) literal.

Like many other solvers, Chaff [22] supports the deletion of added conflict clauses to avoid a memory explosion. Essentially, Chaff uses scheduled lazy clause deletion. When each clause is added, it is analyzed to determine at what point in the future, if any, the clause should be deleted. The metric used is relevance, such that when more than n (where n is typically 100-200) literals in the clause will become unassigned for the first time, the clause will be marked as deleted. The actual memory associated with deleted clauses is recovered with an infrequent monolithic database compaction step.

In Berkmin [14], the authors consider that more recently deduced clauses are more valuable because it took more time to deduce them from the original set of clauses. The learned clause database is considered as a queue (First In First Out). The Berkmin deletion strategy maintains short clauses (size less than 8) combined with queue representation of the learned clauses database.

MiniSat [11] aggressively deletes learned clauses based on an activity heuristic similar to the one for variables. A learnt clause is considered as irrelevant if its activity or its involvement in recent conflict analysis is marginal [11]. The limit on how many learned clauses are allowed is increased after each restart. This strategy has been improved in MiniSAT 2.2.

More recently, Audemard and Simon [1], use the number of different levels (LBD - Literal Block Distance) involved in a given learnt clause to quantify the quality of learnt clauses. The set of different levels present in a learned clause has been used in [2] to prove the optimality of the First UIP scheme. In [1], clauses with smaller LBD are considered as more relevant. The LBD measure is integrated in MiniSAT solver leading to Glucose, one of the state-of-the-art SAT solvers. LBD based measure is also exploited in Lingeling [7], SAT13 [21] the CDCL-based SAT solver designed by D. Knuth, and some other SAT solvers entering the last SAT competition 2013 (e.g. *gluebit_clasp* 1.0, *BreakIDGlucose* 1, *glueminisat* 2.2.7j). Another dynamic management policy of the learnt clauses database is proposed in [3]. It is based on a dynamic freezing and activation principle of the learnt clauses. At a given search state, using a relevant selection function based on progress saving [23], it activates the most promising learned clauses while freezing irrelevant ones.

In [16], we proposed two new measures to predict the quality of learned clauses. The first one is based on the backtracking level (BTL), while the second is based on a notion of distance, defined as the difference between the maximum and minimum assignment levels of the literals involved in the learned clause.

Finally, size-bounded clause sharing strategies are also considered in several portfolio and divide and conquer based parallel SAT solvers (e.g. ManySAT [17] and PMSat [13]).

3 Technical Background and Preliminary Definitions

A propositional formula $\mathcal{F}$ in *Conjunctive Normal Form (CNF)* is a conjunction of *clauses*, where a clause is a disjunction of *literals*. A literal is a positive (x) or negated ($\neg x$) propositional variable. The two literals x and $\neg x$ are called *complementary*. Let c be a clause, $|c|$ denotes the size of c (its number of literals). A *unit* clause is a clause containing only one literal (called *unit literal*), while a binary clause contains exactly two literals. An *empty clause*, denoted $\bot$, is interpreted as false (unsatisfiable), whereas an *empty CNF formula*, denoted $\top$, is interpreted as true (satisfiable). The set of variables occurring in $\mathcal{F}$ is denoted $V_{\mathcal{F}}$. An *interpretation* ρ of a propositional formula $\mathcal{F}$ is a function which associates a value $\rho(x) \in \{false, true\}$ to some of the variables $x \in V_{\mathcal{F}}$. ρ is *complete* if it assigns a value to every $x \in V_{\mathcal{F}}$, and *partial* otherwise. An interpretation ρ is alternatively represented by a set of literals, i.e., $\rho = \bigcup_{x \in V_{\mathcal{F}}} f(x)$, where $f(x) = x$ (respectively $f(x) = \neg x$), if $\rho(x) = true$ (respectively $\rho(x) = false$). A *model* of a formula $\mathcal{F}$ is an interpretation ρ that satisfies the formula.

Let us now introduce some notations and terminology on SAT solvers based on the Davis Logemann Loveland procedure, commonly called DPLL [9]. DPLL is a backtrack search procedure; at each node the assigned literals (decision literal and the propagated ones) are labeled with the same *decision level* starting from 1 and increased at each branching. The current decision level is the highest decision level in the assignment stack. When a conflict is encountered After backtracking, some variables are unassigned, and the current decision level is decreased accordingly. At level i, the current partial assignment ρ can be represented as a sequence of decision-propagation of the form $\langle (x_k^i), x_{k_1}^i, x_{k_2}^i, \ldots, x_{k_{n_k}}^i \rangle$ where the first literal x_k^i corresponds to the decision literal x_k assigned at level i and $x_{k_j}^i$ for $1 \leqslant j \leqslant n_k$ represent (unit) literals propagated at level i. Such a partial interpretation (sequence of decision-propagations) associated to a given node of the search tree is called a *partial ordered interpretation*. Let $x \in \rho$, we denote by $lev(x, \rho)$ the assignment level of x in ρ. Let ρ be a partial interpretation, and c a clause, we define c^i as the projection of c to the literals of c assigned at level i, i.e., $c^i = \{x | lev(x, \rho) = i\}$. This set is called a block in [1].

A conflict driven clause learning (CDCL) SAT solver explore the search space by making successively a sequence of decision/propagation. When a conflict is encountered a conflict clause is derived by resolution on the clauses involved in the unit propagation process (encoded as an implication graph). Such learned conflict clause is then added to the learned clauses database. It allows us to produce a unit (asserting) literal at an earlier level. Then, the CDCL-based SAT solver backtracks to that level, propagates the asserting literal and repeats the sequence of decision/propagation until a model is found or an empty clause is derived. The search regularly restarts, and the learned clauses database is regularly reduced by eliminating irrelevant clauses. These different components are interrelated. For example, the restart have strong effect on the learning component [18,8,25,24,17,4]. Some of the state-of-the-art SAT solvers such as *Glucose* use aggressive restart, very helpful for solving unsatisfiable SAT instances.

4 Size, Random and FIFO Based Deletion Strategies

As mentioned in the introduction, our main aim is to first quantify the performance gap between the state of the art learned clause deletion strategies and some basic strategies including random deletion based one. In this section, we illustrate the performance of size, random and FIFO deletion strategies. All these static strategies are integrated without any other modification to MiniSAT 2.2. The three new versions of MiniSAT are obtained as follows:

- *Size-MiniSAT*: at each conflict, the activity of the learned clause c is set to $|c|$, i.e,. $\mathcal{A}(c) = |c|$
- *Rand-MiniSAT*: at each conflict, the activity of the learned clause c is set to random real value $w \in [0..1]$, i.e., $\mathcal{A}(c) = drand(random_seed)$. We used exactly the random function of MiniSAT with the same random_seed to allow reproducible results.
- *FIFO-MiniSAT*: In this version, the learned clauses are managed using a queue. Each time a reduction is performed, the oldest clauses are deleted.

Solvers	#Solved (#SAT - #UNSAT)	Average Time
MiniSAT	201 (122 - 79)	956.78s
Glucose	216 (104 - 112)	**807.62s**
Lingeling	**233** (119 - **114)**	1090.99s
Size-MiniSAT	220 (**126** - 94)	1226.27s
Rand-MiniSAT	191 (121 - 70)	1071.05s
FIFO-MiniSAT	173 (119 - 54)	862.50s

Table 1. A Comparative Evaluation of Size/Rand/FIFO-MiniSAT

For all the deletion strategies defined in this paper, clauses with smaller activities are considered more relevant. For the deletion frequency and the amount of deleted clauses, we follow the same strategy implemented in MiniSAT 2.2. For all the experiments presented in this paper, we run the SAT solvers on the 300 application instances taken from SAT competition 2013. All the instances are preprocessed by SatElite[10] before running the SAT solver. We used Intel Xeon quad-core machines with 32GB of RAM running at 2.66 Ghz. For each instance, we used a timeout of 5000 seconds of CPU time.

In Table 1, we give the comparative experimental evaluation of our three version of *MiniSAT* with *Glucose* 3.0 and *Lingeling* 2013 the best solvers of the SAT competition 2013 (application category). In the second column, we give the total number of solved instances (#Solved). We also mention, the number of instances proven satisfiable (#SAT) and unsatisfiable (#UNSAT) in parenthesis. The third column show the average CPU time in seconds (total time on solved instances divided by the number of solved instances).

From this first experiment, the state-of-the-art SAT solvers *Lingeling* 2013 and *Glucose* 3.0 solve 233 and 216 instances respectively. The solver $MiniSAT$ solve 201 instances. Let us comment on the performance of the other versions of MiniSAT without any other enhancement. The $Size\text{-}MiniSAT$ solver is able to solve 220 instances (4 more instance than *Glucose* 3.0). Looking at the average CPU time, Glucose 3.0 performs better. Other important observation can be drawn from this first experiment. On satisfiable instances $Size\text{-}MiniSAT$ performs better than all the other solvers. It solves 126 instances. The last remark that can be drawn, $Rand\text{-}MiniSAT$ solves more satisfiable instances (121) than the state-of-the-art solvers *Glucose* 3.0 and *Lingeling* 2013. Finally $FIFO\text{-}MiniSAT$ is clearly the worst, particularly on unsatisfiable instances (only 54 instances are solved). However, it remains competitive on satisfiable instances and solves 15 instances more than *Glucose* 3.0.

5 Size-Bounded Randomized Strategy

The remarquable results obtained in Section 4 by $Size\text{-}MiniSAT$, where short clauses are considered relevant is clearly encouraging. In this section, we propose a new simple reduction strategy, called Size-Bounded Randomized strategy (in short SBR), that combines maintaing short clauses (of size bounded by k), while deleting randomly clauses of size greater than k. One of the main drawback of the clause-size based activity, is that larger clauses are often considered irrelevant. On some hard problems, such drastic selection of learned short clauses might have negative effects. To overcome this problem, we introduced some randomization to the size based approach. More precisely, SBR strategy is defined as follows: given a upper bound k on the length of the learned clauses, each time a learned clause c is derived, if its size is less than k then $\mathcal{A}(c) = |c|$, otherwise $\mathcal{A}(c) = k + drand(random_seed)$. In other words, we still prefer short clauses, while for clauses of size larger than k, they are considered of equal size k with an additional random value. In this way, larger clauses can be selected randomly. The solver $SBR(k)\text{-}MiniSAT$ (where k is the upper bound size) implements the above described strategy. To determine the best upper bound size k, we run $SBR(k)\text{-}MiniSAT$ with $k = 2$, 5, 10, 12 and 15.

The results are depicted in Table 2. $SBR(12)\text{-}MiniSAT$ obtains the best performances. It solves 236 instances, 3 instance more than *Lingeling* 2013 and 20 instances more than *Glucose* 3.0. These remarquable results, bring up to date, old strategies such as size-bounded strategies proposed more than fifteenth years ago [27,5,6]. The results also show that adding some randomization to allow selecting clauses of larger size allows us to introduce some diversification to the resolution derivation of CDCL-Based SAT solvers. It is important to note that our proposed implementation can be obtained by adding three lines of codes to $MiniSAT$ 2.2 without any tuning or additional enhancements. It is worth noticing that the solvers $SBR(k)\text{-}MiniSAT$, for $k = 5, 10, 12, 15$, outperform the solver *Glucose* 3.0.

Solvers	#Solved (#SAT - #UNSAT)	Average Time
Glucose 3.0	216 (104 - 112)	**807.62s**
Lingeling 2013	233 (119 - **114)**	1090.99s
$SBR(2)$-$MiniSAT$	196 (122 - 74)	1064.17s
$SBR(5)$-$MiniSAT$	218 (118 - 100)	1213.36s
$SBR(10)$-$MiniSAT$	231 (124 - 107)	1158.25s
$SBR(12)$-$MiniSAT$	**236** (**128** - 108)	1266.53s
$SBR(15)$-$MiniSAT$	228 (120 - 108)	1226.96s

Table 2. A Comparative Evaluation of $SBR(k)$-$MiniSAT$

6 Towards Relevant-Based Deletion Strategies

Reinforced by the interest of keeping short clauses, our goal in this section is to show if some other variants of these size-based strategies, can lead to better performance. More precisely, our aim is to design dynamic deletions strategies, where the activities of the learned clauses are updated during search. In most of the state-of-the-art SAT solvers, the activities of the learned clauses are updated dynamically.

6.1 Clause Relevance-Based Dynamic Measures

Let us explain the idea that has guided us in the design of these new dynamic variants. Intuitively, *a short learned clause with literals assigned most often at the top of the current branch of the search tree is considered more relevant.*

Given a partial interpretation ρ and c a clause from the learned clause database. Our first dynamic size based deletion strategy is defined as follows: the initial activity of c is set to $|c|$. Suppose now that $l \in \rho$ is assigned by unit propagation at level d, thanks to the learnt clause c. In this case, the new activity of c is set to d if $d < |c|$. This measure allows us to keep learned clauses that are more likely to cut branches at the top of the search tree. For example, for two clauses of equal size, we prefer those containing literals assigned at the top of the search tree. By integrating this measure to MiniSAT 2.2, we obtain our first variant $SizeD$-$MiniSAT$.

Following the same idea, we propose a second version $Size(k)D$-$MiniSAT$. The goal is to introduce a threshold k on the length of the clauses, in order to update only the activity of clauses greater than k, and keeping the activity of short clauses (size less than k) static. More precisely, the initial activity of c is set as follows: if $|c| \leqslant k$ then $\mathcal{A}(c) = |c|$, otherwise $\mathcal{A}(c) = k + |c|$. Each time, a learned clause c is the reason of a unit propagated literal at decision level d, its activity is updated as follows: if $k + d < \mathcal{A}(c)$ then $\mathcal{A}(c) = k + d$. In this way, the activity of a clause with size less than k is not updated.

The third variant is defined as follows. Let ρ be a partial interpretation leading to a conflict at decision level d and c its associated learned clause. The activity of c is initially set to $\mathcal{A}(c) = \sum_{i=1}^{d} i \times |c^i|$. During search, each time c is

the reason of a unit propagated literal, and if its new activity w.r.t. the current interpretation is better, the activity of c is updated. A clause c is considered better than a clause c' if $\mathcal{A}(c) < \mathcal{A}(c')$. This new strategy leads us to the new solver denoted $RelD\text{-}MiniSAT$.

Solvers	#Solved (#SAT - #UNSAT)	Average Time
Glucose 3.0	216 (104 - 112)	**807.62s**
Lingeling 2013	**233** (119 - **114)**	1090.99s
$Size(12)D\text{-}MiniSAT$	**233** (**123** - 110)	1153.9s
$SizeD\text{-}MiniSAT$	231 (120 - 111)	1174.97s
$RelD\text{-}MiniSAT$	222 (122 - 100)	1157.31s

Table 3. A Comparative Evaluation of Dynamic Strategies

In Table 3, we present the comparative results of the three new dynamic relevant based deletions strategies. As we can see, *Lingeling* 2013 and $Size(12)D\text{-}MiniSAT$ present comparable performances (233 solved instances). We also observe that *Lingeling* 2013 is better (respectively worse) than $Size(12)D\text{-}MiniSAT$ on unsatisfiable (respectively satisfiable) instances. The solver MiniSAT integrating our three relevant based deletion strategies outperforms the solver *Glucose* 2013.

In this section, three dynamic strategies are proposed where the activities of the learned clauses are updated dynamically during search. As a summary, size-bounded randomized strategy $SBR(12)\text{-}MiniSAT$ remains the best, as it solves more instances (236 solved instances) than all the proposed strategies proposed in this paper including the state-of-the-art SAT solvers *Lingeling* 2013 and *Glucose* 3.0.

All the proposed strategies can be easily integrated to MiniSAT using few lines of codes.

6.2 Substituting LBD with Clause Size Inside Glucose

In the previous sections, we compared our deletion strategies integrated in MiniSAT with the state-of-the-art SAT solvers *Glucose* 3.0 and *Lingeling* 2013. These two solvers exploit LBD based measure for managing the learned clauses database. The question that we want to answer is the following: what about substituting LBD measure with clause size inside *Glucose* 3.0 and *Lingeling* 2013? As these two solvers exploit LBD based measure for managing the learned clauses database, we only modify *Glucose* 3.0 to evaluate the effect of substituting LBD with dynamic clause size ($Size(k)D$) presented in Section 6.1. The obtained SAT solver is called $Size(k)D$-*Glucose* 3.0, where k is a clause size threshold. This new solver is obtained by substituting LBD with Size in the two following parts of *Glucose* 3.0:

– *Initial learned clause activity*: when a learned clause c is derived, its activity is set as follows: $\mathcal{A}(c) = |c|$ if $|c| < k$, otherwise $\mathcal{A}(c) = k + |c|$.
– *Dynamic update of clause activity*: at each conflict, the activity of the learned clauses involved in the derivation of the asserting clause is updated as follows: Let d be the conflict decision level. if $k + d < \mathcal{A}(c)$ then $\mathcal{A}(c) = k + d$. Obviously, the activity of a clause of size less than k is not updated.

The $SBR(k)$-*Glucose* 3.0 integrates size-bounded randomized deletion strategy. It is obtained as follows: each time a learned clause c is derived, it receives $\mathcal{A}(c) = |c|$ if $|c| \leqslant k$, otherwise $\mathcal{A}(c) = k + irand(random_seed, |V_{\mathcal{F}}|)$, where $irand(random_seed, |V_{\mathcal{F}}|)$ return a number between 0 and $|V_{\mathcal{F}}|$. This activity remains unchanged during search.

Solvers	#Solved (#SAT - #UNSAT)	Average Time
Glucose 3.0	216 (104 - 112)	**807.62s**
Size(12)*D-Glucose*	**224** (103 - **121**)	919.16s
Size(15)*D-Glucose*	219 (100 - 119)	944.36s
SBR(15)*-Glucose*	222 (**105** - 117)	1001.34s

Table 4. A Comparative Evaluation of *Glucose*, *Size*(k)*D-Glucose* and $SBR(k)$*-Glucose*

This experiment demonstrates that the clause size is clearly more relevant than LBD/glue measure. As we can see, from the result depicted in Table 4, *Size*(12)*D-Glucose* 3.0 solves 8 instances more than *Glucose* 3.0. More interestingly, on unsatisfiable instances, our version solves 9 instances more than *Glucose* 3.0. On satisfiable instances, both solvers present similar behavior. *Glucose* 3.0 solves only one additional instance. This experiment give us a definitive illustration that clause size based measure is better than LBD.

7 Discussion

As a summary, we demonstrate that the size of the clause remains the best metric to quantify the usefulness of a given clause. The lesson that can be drawn from this study, is that predicting the best clauses to be maintained during search deserve further investigation. In our opinion, the performance of the LBD measure can be explained by the fact that it is really related to the size of the clauses. Indeed, we have $2 \leqslant LBD(c) \leqslant |c|$. The clause size is an upper bound of the LBD measure. For instance, the LBD of binary clauses is 2. Hence, the strategy defined in Glucose favors in some sens maintaining short clauses. In [1], the authors mention that "*A good learning schema should add explicit links between independent blocks of propagated (or decision) literals. If the solver stays in the same search space, such a clause will probably help reducing the number of next decision levels in the remaining computation*". The intuition given by the

authors, to understand the importance of clauses of LBD 2: "*the variable from the last decision level will be glued with the block of literals propagated above. We suspect all those clauses to be very important during search*". Let us discuss these claims. First any learned clause add explicit links between independent blocks of propagated literals. Secondly, any learned clause also help to prune the search space and then to reduce the number of decision. The question we ask is the following: let $c_1 = (x_1 \vee x_2 \vee x_3)$ and $c_2 = (y_1 \vee y_2 \vee, \ldots, \vee y_n)$ with $n > 3$, $LBD(c_1) = 3$ and $LBD(c_2) = 2$, which one is relevant? As you can guess, our preference leans to the first clause.

An important remark, all the strategies presented in this paper can be integrated easily (few lines of code) to any CDCL-based SAT solver. A dedicated web page including the proposed learned clauses database deletion strategies integrated to MiniSAT is currently under construction (http://www.cril.fr/∼sais).

8 Conclusion and Futur Works

In this paper, we addressed the learned clauses database management problem. We demonstrated that size-bounded learning strategies proposed more than fifteenth years ago [27,5,6] remains a good measure to predict the quality of learned clauses. We have shown that adding randomization to size bounded learning is a nice way to achieve controlled diversification. It allows us to favor short clauses, while maintaining a small fraction of large clauses necessary for deriving resolution proofs on some SAT instances. The experimental results, shows that Size-Bounded Randomized strategy integrated to MiniSAT can achieve better performance than the state-of-the-art SAT solvers. Convinced by the importance of short clauses, we proposed several efficient dynamic variant that aims to maintain short clauses containing literals that occurs more often at the top of the search tree. Our last evaluation shows that substituting LBD with clause size based measure inside Glucose leads to better performance on both SAT and UNSAT instances. This paper, open several perspectives including the refinement of the strategies proposed in this paper.

References

1. G. Audemard and L. Simon. Predicting learnt clauses quality in modern sat solvers. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence, IJCAI'09*, pages 399–404, 2009.
2. Gilles Audemard, Lucas Bordeaux, Youssef Hamadi, Said Jabbour, and Lakhdar Sais. A generalized framework for conflict analysis. Technical Report MSR-TR-2008-34, Microsoft Research, Feb 2008.
3. Gilles Audemard, Jean-Marie Lagniez, Bertrand Mazure, and Lakhdar Sais. On freezing and reactivating learnt clauses. In *Proceedings of the 14th International Conference on Theory and Applications of Satisfiability Testing (SAT'2011)*, pages 188–200, 2011.

4. Gilles Audemard and Laurent Simon. Refining restarts strategies for sat and unsat. In *18thInternational Conference on Principles and Practice of Constraint Programming(CP'2012)*, pages 118–126, 2012.
5. Roberto J. Bayardo and Daniel P. Miranker. A complexity analysis of space-bounded learning algorithms for the constraint satisfaction problem. In *In Proceedings of the Thirteenth National Conference on Artificial Intelligence*, pages 298–304, 1996.
6. Roberto J. Bayardo, Jr. and Robert C. Schrag. Using CSP look-back techniques to solve real-world SAT instances. In *AAAI*, pages 203–208, 1997.
7. A. Biere. Lingeling and friends entering the sat challenge 2012. In A. Balint, A. Belov, A. Diepold, S. Gerber, M. Jarvisalo, , and C. Sinz (editors), editors, *Proceedings of SAT Challenge 2012: Solver and Benchmark Descriptions*, pages 33–34, University of Helsinki, 2012. vol. B-2012-2 of Department of Computer Science Series of Publications B.
8. Armin Biere. Adaptive restart strategies for conflict driven sat solvers. In *Proceedings of the 11th International Conference on Theory and Applications of Satisfiability Testing (SAT'08)*, pages 28–33, 2008.
9. M. Davis, G. Logemann, and D. W. Loveland. A machine program for theorem-proving. *Communications of the ACM*, 5(7):394–397, 1962.
10. Niklas Eén and Armin Biere. Effective preprocessing in sat through variable and clause elimination. In *Proc. of SAT'05*, pages 61–75, 2005.
11. Niklas Eén and Niklas Sörensson. An extensible sat-solver. In *Proceedings of the Sixth International Conference on Theory and Applications of Satisfiability Testing (SAT'03)*, pages 502–518, 2003.
12. David Gelperin. Deletion-directed search in resolution-based proof procedures. In *Proceedings of the 3rd International Joint Conference on Artificial Intelligence (IJCAI'73)*, pages 47–50, San Francisco, CA, USA, 1973. Morgan Kaufmann Publishers Inc.
13. Luís Gil, Paulo Flores, and Luís M. Silveira. PMSat: a parallel version of MiniSAT. *Journal on Satisfiability, Boolean Modeling and Computation*, 6:71–98, 2008.
14. Eugene Goldberg and Yakov Novikov. Berkmin: A fast and robust sat-solver. *Discrete Applied Mathematics*, 155(12):1549 – 1561, 2007.
15. Carla P. Gomes, Bart Selman, and Henry A. Kautz. Boosting combinatorial search through randomization. In *AAAI/IAAI*, pages 431–437, 1998.
16. Long Guo, Saïd Jabbour, and Lakhdar Saïs. Stratégies d'élimination des clauses apprises dans les solveurs sat modernes. In *9ièmes Journées Francophones de Programmation par Contraintes (JFPC'13)*, pages 147–156, Aix en provence, France, 2013.
17. Youssef Hamadi, Sad Jabbour, and Lakhdar Sas. Manysat: a parallel sat solver. *Journal on Satisfiability, Boolean Modeling and Computation (JSAT)*, 6:245–262, 2009.
18. Jinbo Huang. The effect of restarts on the efficiency of clause learning. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI'07)*, pages 2318–2323, 2007.
19. Hadi Katebi, KaremA. Sakallah, and Joao P. Marques-Silva. Empirical study of the anatomy of modern sat solvers. In KaremA. Sakallah and Laurent Simon, editors, *Theory and Applications of Satisfiability Testing - SAT 2011*, volume 6695 of *Lecture Notes in Computer Science*, pages 343–356, 2011.
20. Henry A. Kautz, Eric Horvitz, Yongshao Ruan, Carla P. Gomes, and Bart Selman. Dynamic restart policies. In *The Eighteenth National Conference on Artificial Intelligence (AAAI'02)*, pages 674–681, 2002.

21. D. Knuth. Sat13. http://www-cs-faculty.stanford.edu/∼uno/programs/sat13.w.
22. Matthew W. Moskewicz, Conor F. Madigan, Ying Zhao, Lintao Zhang, and Sharad Malik. Chaff: Engineering an efficient sat solver. In *38th Design Automation Conference (DAC'01)*, pages 530–535, 2001.
23. Knot Pipatsrisawat and Adnan Darwiche. A lightweight component caching scheme for satisfiability solvers. In *SAT*, pages 294–299, 2007.
24. Knot Pipatsrisawat and Adnan Darwiche. On the power of clause-learning sat solvers with restarts. In *15th International Conference on Principles and Practice of Constraint Programming (CP'09)*, pages 654–668, 2009.
25. Knot Pipatsrisawat and Adnan Darwiche. Width-based restart policies for clause-learning satisfiability solvers. In *Proceedings of the 12th International Conference on Theory and Applications of Satisfiability Testing (SAT'09)*, pages 341–355, 2009.
26. Knot Pipatsrisawat and Adnan Darwiche. On modern clause-learning satisfiability solvers. *Journal of Automated Reasoning*, 44(3):277–301, 2010.
27. João P. Marques Silva and Karem A. Sakallah. Grasp - a new search algorithm for satisfiability. In *International Conference on Computer-Aided Design (ICCAD'96)*, pages 220–227, 1996.